arXiv:1401.6694v2 [cs.SC] 2 May 2014

Multivariate sparse interpolation using randomized Kronecker substitutions

Andrew Arnold
Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
a4arnold@uwaterloo.ca

Daniel S. Roche
Computer Science Department
United States Naval Academy
Annapolis, Maryland, USA
roche@usna.edu

May 5, 2014

Abstract

We present new techniques for reducing a multivariate sparse polynomial to a univariate polynomial. The reduction works similarly to the classical and widely-used Kronecker substitution, except that we choose the degrees randomly based on the number of nonzero terms in the multivariate polynomial. The resulting univariate polynomial often has a significantly lower degree than the Kronecker substitution polynomial, at the expense of a small number of term collisions. As an application, we give a new algorithm for multivariate interpolation which uses these new techniques along with any existing univariate interpolation algorithm.

1 Introduction

We consider the problem of determining the coefficients and exponents of an unknown sparse multivariate polynomial, given a "black box" procedure that allows for its evaluation at any chosen point.

Our new technique is a variant on the classical Kronecker substitution. Say f is an n-variate polynomial with max degree less than D and coefficients in a ring R. The Kronecker substitution produces a univariate polynomial $g \in \mathsf{R}[z]$ by substituting powers of z in the evaluation of f:

$$g(z) = f\left(z, z^D, z^{D^2}, \ldots, z^{D^{n-1}}\right).$$

This map is invertible because there is a one-to-one correspondence between terms in f and in g, but the price of such convenience is an exponential increase in the degree.

For example, consider the following bivariate polynomial:

$$f = 3x^9y - 2x^5y^4 - y^6 + x^2y^9.$$

The standard Kronecker substitution would be

$$g = f(z, z^{10}) = 3z^{19} - 2z^{45} - z^{60} + z^{92}.$$

Every term in g comes from a single term in f, and the original exponents can be determined by the base-D expansion of exponents in g.

Our randomized Kronecker substitution is also a map from multivariate to univariate polynomials obtained by evaluating at powers of a single indeterminate. We choose n integers $(s_1, \ldots, s_n)$ at random and perform the substitution $g(z) = f(z^{s_1}, \ldots, z^{s_n})$. When these integers are not too large (in particular, if each $s_i < D^{n-1}$), the degree of g will be less than in the usual Kronecker substitution.

The price of a decreased degree is that the map is no longer invertible, for two reasons. First, we may have two or more distinct terms in f converge to a single term in g. This is called a *collision*. The second difficulty is that the original multivariate exponents cannot be determined directly from the terms in a single substitution g. We will show how performing $O(n + \log \#f)$ such random substitutions can overcome both difficulties.

In the example above, choose $s_1 = 5$ and $s_2 = 2$, so that

$$g_{5,2}(z) = f(z^5, z^2) = -z^{12} + z^{28} - 2z^{29} + 3z^{47}.$$

In this case $\#g = \#f$, so there were no term collisions, even though their order has changed. The advantage is that the degree is significantly less than that from the usual Kronecker substitution.

Choosing instead $s_1 = 2$ and $s_2 = 5$, the result is

$$g_{2,5}(z) = f(z^2, z^5) = 3z^{23} - 3z^{30} + z^{49},$$

which again has a reduced degree, but in this case we have a *collision*: The two terms $-2x^5y^4 - y^6$ in f collided to produce a single term $-3z^{30}$ under this substitution.

Nonetheless, for the two terms not involved in a collision, both images can be used to recover the original terms in f. The two terms with coefficient 3 in the substitutions are $3z^{47}$ and $3z^{23}$. This produces the linear system

$$\begin{bmatrix} 5 & 2 \\ 2 & 5 \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} 47 \\ 23 \end{bmatrix},$$

which is solved to reveal exponents $u = 9$ and $v = 1$ of the original term $3x^9y$.

Crucial to our success is determining a suitable set of integers from which to choose the s_i. Section 3 provides bounds to construct such sets with provably many "good" choices which will not produce many collisions. Recovering the original terms requires a way to correlate like terms in the images g, and Section 4 shows how a separate randomization allows the coefficients to be used to identify like terms in separate images g. Section 5 describes a multivariate interpolation algorithm that makes use of these randomizations.

2 Related work

Polynomial interpolation dates from the 18th century, where the goal was to discover a polynomial approximation to an unknown function by collecting sufficiently many observations. It is often beneficial to search for a *sparse* model, where only a bounded number of terms in the polynomial are nonzero (see, e.g., Candés, Romberg, and Tao [2006]).

Our setting is more restrictive in two senses. First, we require the ability to choose the points at which the unknown polynomial is evaluated. In this *black-box* interpolation setting, the polynomial a procedure which can be probed with any desired input, at some cost (see, e.g., Grigoriev, Karpinski, and Singer [1990], Mansour [1995], Giesbrecht, Labahn, and Lee [2009]).

We also require that a sparse polynomial truly exists "inside the box", and all evaluations — even those with numerical noise — come from the same T-sparse polynomial. This is different from the general setting above where the true function need not be T-sparse or indeed a polynomial at all. (Allowing for outliers as in Comer, Kaltofen, and Pernet [2012] provides another option.)

Algorithmic progress in sparse interpolation can be divided into two categories depending on the complexity. Algorithms with polynomial dependence on the partial degree bounds, sparsity bound, and number of variables are called *sparse* algorithms [Zippel, 1979, 1990, Huang and Rao, 1999]. Those with only logarithmic dependence on the degree are *supersparse* and can be useful even in the case of univariate polynomials Ben-Or and Tiwari [1988], Avendaño, Krick, and Pacetti [2006].

As shown in the introduction, Kronecker subtitution [Kronecker, 1882] can turn a univariate interpolation into a multivariate one, by evaluating the multivariate polynomial at high powers of the univariate interpolation points. This is especially effective in conjunction with supersparse algorithms [Kaltofen, 2010].

Zippel's sparse interpolation method proceeds one variable at a time and uses randomization to identify which portions of the unknown sparse polynomial vanish. This is another way to turn any univariate algorithm into a multivariate one Kaltofen and Lee [2003].

There is recent work on implementations of sparse interpolation Javadi and Monagan [2010], van der Hoeven and Lecerf [2014] and applications to GCD and factorization [Javadi and Monagan, 2009, Berthomieu and Lecerf, 2012]. The Kronecker substitution in particular has been applied to integer and polynomial multiplication [Schönhage, 1982, Harvey, 2009].

Sparse interpolation has a strong connection to the important theoretical problem of polynomial identity testing (see Shpilka and Yehudayoff [2010] for a survey). Of particular interest is Klivans and Spielman [2001], where a similar technique to ours is used for identity testing and interpolation. A significant advantage of their algorithms is the very small number of random bits required for the computation.

Our randomized Kronecker substitution provides yet another way to transform any univariate interpolation algorithm into a multivariate one. Table 1 summarizes our contribution as compared to Klivans and Spielman [2001], Zip-

pel's method, and the normal Kronecker substitution. The table shows, for each method, how many times a univariate interpolation algorithm must be called, and the degree of the univariate polynomials that must be interpolated, relative to a given bound D on the max degree of the unknown multivariate polynomial, and the number of variables n. (The sparsity bound for the univariate algorithms will always be simply T.)

Table 1: Uni-to-multivariate interpolation methods

	# of reductions	Degree
Kronecker substitution	1	D^n
Zippel	nT	D
Klivans & Spielman	n	$O(n^2T^2D)$
This paper ($n = 2$ case)	$O(\log T)$	$O\left(\sqrt{T}D \log D\right)$
This paper ($n \geqslant 3$ case)	$O(n + \log T)$	$O\left(TD\right)$

The total cost of a multivariate interpolation algorithm obtained this way depends on the choice of univariate interpolation algorithm. Writing D_1 for the univariate degree (the second column of Table 1), there are two options: dense methods which require D_1 black box probes and $\widetilde{O}\left(D_1\right)$ field operations, or supersparse methods based on Ben-Or and Tiwari which require only $O(T)$ probes and $\widetilde{O}\left(T \log^2 D_1\right)$ field operations. In Table 2 we compare two standard approaches against our own: Kronecker substitution with a supersparse univariate algorithm, and Zippel's method with a dense univariate algorithm.

Table 2: Overall multivariate interpolation

reduction	univariate alg.	# of probes	# of field ops
Kronecker	supersparse	$2T$	$\widetilde{O}\left(n^2T \log^2 D\right)$
Ours	supersparse	$\widetilde{O}\left(nT\right)$	$\widetilde{O}\left(nT \log^2 D\right)$
Zippel	dense	nTD	$\widetilde{O}\left(nTD\right)$
Ours ($n = 2$)	dense	$\widetilde{O}\left(\sqrt{T}D\right)$	$\widetilde{O}\left(\sqrt{T}D\right)$
Ours ($n \geqslant 3$)	dense	$\widetilde{O}\left(nTD\right)$	$\widetilde{O}\left(nTD\right)$

3 Randomized substitutions

For an unknown polynomial $f \in \mathsf{R}[x_1, \ldots, x_n]$, our main technical contribution is a way of choosing integers $s_1, \ldots, s_n \in \mathbb{Z}$ such that the substitution $g = f\left(z^{s_1}, \ldots, z^{s_n}\right)$ results in a lower degree than the usual Kronecker substitution, while probably not introducing too many term collisions.

3.1 Bivariate substitutions

We begin with the case of $n = 2$ variables, where our result is stronger than the general case and we always choose the random substitution exponents s_1, s_2 to be prime numbers. Bivariate polynomials naturally constitute a large portion of multivariate polynomials of interest, and they correspond to the important case of converting between polynomials in $\mathbb{Z}[x]$ and multiple-precision integers [Schönhage, 1982, Harvey, 2009].

Throughout this subsection, we assume $f \in \mathsf{R}[x, y]$ is an unknown bivariate polynomial, written as

$$f = a_1 x^{u_1} y^{v_1} + a_2 x^{u_2} y^{v_2} + \cdots + a_t x^{u_t} y^{v_t}. \tag{3.1}$$

We further assume upper bounds D_x, D_y on the $\deg_x f$ and $\deg_y f$, respectively, and $T \geqslant t$ on the number of nonzero terms $\#f$.

The general idea here is to perform the substitution

$$g(z) = f(z^p, z^q) \tag{3.2}$$

for random chosen prime numbers p and q. We want to choose p and q as small as possible, so as to minimize $\deg g$, but large enough so that there are not too many collisions.

Our approach to choosing primes is based on the following technical lemma, which shows how to guarantee a high probability success while minimizing the degree of g.

Lemma 3.1. *Let $f \in \mathsf{R}[x, y]$ with partial degrees less than D_x, D_y and at most than T nonzero terms, $0 < \mu < 1$ be a chosen bound on the probability of failure, and $1 \leqslant i \leqslant T$ be the index of a nonzero term in f. Define*

$$B = \frac{25\,(T-1) \ln D_x \ln D_y}{9\mu},$$

$$\lambda_p = \max\left(20.5, \sqrt{\frac{BD_y}{D_x}}\right) \quad \textit{and} \quad \lambda_q = \max\left(20.5, \sqrt{\frac{BD_x}{D_y}}\right)$$

By choosing primes p, q uniformly at random from the ranges $[\lambda_p, 2\lambda_p]$ and $[\lambda_q, 2\lambda_q]$ respectively, the probability that the ith term of f collides in $f(z^p, z^q)$ is less than μ.

Proof. The primes p and q constitute the random choices in this discussion. We say a pair (p, q) is *bad* if the term $a_i x^{u_i} y^{v_i}$ collides with any other term in $f(z^p, z^q)$. We will show that the number of bad pairs (p, q) is at most μ times the total number of prime pairs that could be chosen.

We begin with a simple lower bound on the latter quantity. Applying equation (3.8) in Rosser and Schoenfeld [1962] twice shows that the total number of ordered pairs (p, q) is at least

$$\frac{9\lambda_p \lambda_q}{25 \ln \lambda_p \ln \lambda_q}. \tag{3.3}$$

Next we obtain an upper bound on the number of bad pairs by counting the total number of times the ith term collides in every possible $f(z^p, z^q)$.

Observe that for any set S of nonzero integers, and any bound λ, the number of times any prime $p \geqslant \lambda$ divides any of the integers in S is at most

$$\frac{|S| \ln \max(S)}{\ln \lambda}. \tag{3.4}$$

A collision with the ith term occurs whenever $u_i p + v_i q = u_j p + v_j q$ for some $j \neq i$, $1 \leqslant j \leqslant T$. This happens only when $p \mid (v_i - v_j)$ and $q \mid (u_i - u_j)$. Since $i \neq j$, these exponent differences $(u_i - u_j)$ and $(v_i - v_j)$ cannot both be zero. Furthermore, if one exponent difference is zero, then the collision can never occur.

Therefore all collisions occur at indices in the set

$$J = \{j \mid 1 \leqslant j \leqslant T \text{ and } (u_i - u_j)(v_i - v_j) \neq 0\}.$$

The total number of times the ith term collides in any $f(z^p, z^q)$ is equal to the sum over all $j \in J$ of the number of pairs (p, q) such that $p \mid (v_i - v_j)$ and $q \mid (u_i - u_j)$.

Now define, for each p, the subset of *possible* collision indices as $J_p = \{j \in J \mid p \text{ divides } (v_i - v_j)\}$.

As each $(v_i - v_j) < D_y$, we have from (3.4) that

$$\sum_{p \geqslant \lambda_p} \#J_p \leqslant (T-1) \ln D_y / \ln \lambda_p.$$

For each prime p, the total number of times the ith term collides is the number of indices $j \in J_p$ such that $q \mid (u_i - u_j)$. As each of these differences is less than D_x, using (3.4) again this sum is at most $(\#J_p \cdot \ln D_x)/(\ln \lambda_q)$.

Therefore the total number of times the ith term collides is at most

$$\sum_{p \geqslant \lambda_p} \frac{\#J_p \ln D_x}{\ln \lambda_q} \leqslant \frac{(T-1) \ln D_x \ln D_y}{\ln \lambda_p \ln \lambda_q}.$$

Using the definition of B and the observation that $\lambda_p \lambda_q = B$, we can rewrite this bound as

$$\frac{9\mu}{25 \ln \lambda_p \ln \lambda_q} B = \mu \frac{9 \lambda_p \lambda_q}{25 \ln \lambda_p \ln \lambda_q},$$

which is exactly μ times (3.3). Hence the probability of choosing a bad pair (p, q) is at most μ. □

Choosing primes p, q from such sets provides a good bound on the degree of the resulting polynomial g.

Corollary 3.2. *Let $f \in \mathsf{R}[x, y]$ with partial degrees less than D_x, D_y and at most T nonzero terms.*

Then for any constant error probability $0 < \mu < 1$, and primes p, q chosen randomly as in Lemma 3.1, the substitution polynomial $g(z) = f(z^p, z^q)$ has degree at most

$$O(\sqrt{T}\sqrt{D_x D_y}\log(D_x D_y)).$$

The degree of a standard Kronecker substitution is $D_x D_y$. Because T is always less than this, the randomized substitution will never have degree more than a logarithmic factor larger than $D_x D_y$. The benefit comes when the polynomial is sparse, i.e., $T \ll D_x D_y$, in which case the randomized substitution has much smaller degree, albeit at the expense of a few collisions.

3.2 Multivariate substitutions

When f has at least 3 variables, the preceding analysis no longer applies. The new difficulty is that potentially-colliding terms could have exponents that differ in two *or more* variables, meaning that the simple divisibility conditions are no longer sufficient to identify every possible collision. Consequently, our randomly-chosen exponents in this case will be somewhat larger, and not necessarily prime.

For this subsection, $f \in \mathsf{R}[x_1, \ldots, x_n]$ is an unknown n-variate polynomial, written as

$$f = a_1\mathbf{x}^{\mathbf{e}_1} + a_2\mathbf{x}^{\mathbf{e}_2} + \cdots + a_t\mathbf{x}^{\mathbf{e}_t}, \tag{3.5}$$

where $\mathbf{x} = (x_1, \ldots, x_n)$ and each $\mathbf{e}_i \in \mathbb{Z}^n$. D and T are upper bounds on the max degree and number of terms in f.

Our general approach here is to choose a random vector $\mathbf{s} = (s_1, s_2, \ldots, s_n)$ of integers below a certain bound, and then perform the substitution $g(z) = f(z^{s_1}, \ldots, z^{s_n})$.

The following lemma, similar in purpose to Lemma 3.1, shows how large the integers in $\mathbf{s}$ must be in order to guarantee a small likelihood of collisions.

Lemma 3.3. *Let $f \in \mathsf{R}[x_1, \ldots, x_n]$ with max degree less than D and at most T nonzero terms, $0 < \mu < 1$ be a chosen bound on the failure probability, and $1 \leqslant i \leqslant T$ be the index of a nonzero term in f. Define λ to be the least prime number satisfying $\lambda \geqslant T/\mu$.*

If integers $s_1, \ldots, s_n$ are chosen uniformly at random from $[0, \lambda - 1]$, then the probability that the ith term of f collides in $f(z^{s_1}, \ldots, z^{s_n})$ is less than μ.

Proof. Adopt the notation $\mathbb{F}_\lambda$ for the finite field with λ elements, which we will represent as $\mathbb{Z}/\lambda\mathbb{Z}$. Write $\mathbf{s} = (s_1, \ldots, s_n)$ for the randomly-chosen vector in $\mathbb{F}_\lambda^n$.

Now let $1 \leqslant j \leqslant T$, $j \neq i$, and consider the jth term of f. Writing $\mathbf{e}_i, \mathbf{e}_j$ for the exponent vectors of these terms as in (3.5), define $\mathbf{d}_j = \mathbf{e}_i - \mathbf{e}_j$, which cannot be the zero vector as i and j are distinct terms. We see that these two terms collide in the substitution $f(z^{s_1}, \ldots, z^{s_n})$ if and only if $\mathbf{d}_j \cdot \mathbf{s} = 0$.

Define $\ell \geqslant 0$ such that λ^ℓ is the largest power of λ that divides every entry in $\mathbf{d}_j$, and write $\mathbf{d}'_j = \mathbf{d}_j/\lambda^\ell$. This means that $\mathbf{d}'_j \in \mathbb{Z}^n$ and $\mathbf{d}'_j \bmod \lambda \neq \mathbf{0}$. Furthermore, $\mathbf{d}_j \cdot \mathbf{s} = 0$ if and only if $\mathbf{d}'_j \cdot \mathbf{s} = 0$.

Now, if $\mathbf{d}'_j \cdot \mathbf{s} = 0$, then this also holds modulo λ, so $\mathbf{s}$ must lie in the $(n-1)$-dimension null space of $\mathbf{d}'_j \bmod \lambda$, call it W, where $W \subset \mathbb{F}_\lambda^n$. The probability

that $\mathbf{s} \in W$ is $1/\lambda$, and therefore the probability that terms i and j collide is at most $1/\lambda$ as well.

Since there are $T-1$ terms that the ith term could collide with, i.e., $T-1$ choices for j, the probability that the ith term collides with any other term is at most $(T-1)/\lambda$. From the definition of λ, this is less than μ. □

Due to Bertrand's postulate, we will have $\lambda < 2T/\mu$. The following corollary shows how this bound on the size of entries in the randomly-chosen $\mathbf{s}$ affects the degree of the reduced univariate polynomial. Compared to D^n, the degree of the univariate polynomial from the usual Kronecker substitution, we see a significant reduction when $T \ll D^{n-1}$.

Corollary 3.4. *Let $f \in \mathsf{R}[x_1, \ldots, x_n]$ with max degree less than D and at most T nonzero terms.*

For any constant error probability $0 < \mu < 1$, and integers $s_1, \ldots, s_n$ chosen randomly as in Lemma 3.3, the polynomial $g(z) = f\left(z^{s_1}, z^{s_2}, \ldots, z^{s_n}\right)$ has degree at most $O(TD)$.

4 Multivariate diversification

Consider $f \in \mathsf{R}[x_1, \ldots, x_n]$ as in (3.5):

$$f = a_1\mathbf{x}^{\mathbf{e}_1} + a_2\mathbf{x}^{\mathbf{e}_2} + \cdots + a_t\mathbf{x}^{\mathbf{e}_t}.$$

Each choice of $\mathbf{s} = (s_1, \ldots, s_n)$ for a randomized Kronecker substitution maps f to a univariate polynomial $g(z) = f(z^{s_1}, \ldots, z^{s_n})$. In order to recover the exponent tuples of the original polynomial f, it will be necessary in the next section to perform multiple such substitutions and correlate terms in each g that correspond to the same unknown term in f.

The notion of *diversification*, introduced by Giesbrecht and Roche [2011], will be used to correlate terms in the substituted polynomials g. The basic idea is that distinct terms in f will be made, through a randomization, to have distinct coefficients.

In fact, as there will be some small number of collisions in each substituted polynomial g, we require the notion of *generalized diversity* from Arnold, Giesbrecht, and Roche [2014]. The idea is that not only must the terms in f have distinct coefficients, but some small number of *sums* of terms in f must additionally be distinct.

This problem of diversification is to choose α from a suitable set so that, with high probability, α is not a root of any in a set $\mathcal{H}$ of polynomials. In the original notion of diversity, $\mathcal{H}$ simply consists of the set of $\frac{T(T-1)}{2}$ pairwise term differences from f. To achieve generalized diversity, we must also consider polynomials h of the form

$$h = \sum_{i \in S} a_i\mathbf{x}^{\mathbf{e}_i} - \sum_{j \in S'} a_j\mathbf{x}^{\mathbf{e}_j},$$

where S and S' each either comprise a single term or a set of terms appearing in a collision.

Definition 4.1 (Diversifying set)**.** *Let $n \geqslant 1$ and bounds D, m, and μ be given, and let $\mathcal{H} \subseteq \mathsf{R}[x_1, \ldots, x_n]$ be an unknown subset of nonzero polynomials satisfying $\#\mathcal{H} \leqslant m$ and with each $h \in \mathcal{H}$ having max degree less than D.*

*We say $A \subseteq \mathsf{R}$ is a (n, D, m, μ)-*diversifying set *if the probability is less than μ that any evaluation point $\boldsymbol{\alpha}$, with entries chosen at random from A, is a root of* any *of the $h \in \mathcal{H}$. That is,*

$$\Pr_{\boldsymbol{\alpha} \in A^n} \left[h(\boldsymbol{\alpha}) \neq 0 \ \forall h \in \mathcal{H} \right] \geqslant 1 - \mu.$$

From the discussion above, we see that the set $\mathcal{H}$ the differences between any of the $t \leqslant T$ single terms and any of the r sets of collisions, which is at most $\#H < \frac{1}{2}(T + r)^2$.

Lemma 4.2. *Let $f \in \mathsf{R}[x_1, \ldots, x_n]$ with degree less than D and at most T nonzero terms, and $0 < \mu < 1$. If there are at most r collisions in some set of evaluations of f, and $\boldsymbol{\alpha} \in \mathsf{R}^n$ is chosen at random from a set $A \subseteq \mathsf{R}$ that is a $(n, D, \frac{1}{2}(T + r)^2, \mu)$-diversifying set, then, with probability at least $1 - \mu$, every coefficient of $f(\alpha_1 x_1, \ldots, \alpha_n x_n)$ is distinct from every other coefficient and from the coefficients of all r collisions.*

Proof. Follows from the discussion above and the definition of a diversifying set. □

From the definition of diversifying set, a simple application of the Schwartz-Zippel lemma could be used to generate diversifying sets as long as the field R is sufficiently large. Theorems 3.1 and 4.6 in Giesbrecht and Roche [2011] define diversifying sets for large finite fields and fixed-precision complex numbers, respectively, in the univariate case $n = 1$. Our more recent work in Arnold et al. [2014] constructs smaller diversifying sets by choosing *vectors* of substitutions, again for the univariate case. We restate these results in our current notation and refer the reader to the aforementioned results for further details.

Corollary 4.3 (Giesbrecht and Roche [2011], Theorem 3.1)**.** *Let bounds D, m, and μ be given. If q is a prime power satisfying $q \geqslant mD/\mu$, then the set $\mathbb{F}_q^*$ of all nonzero elements in the finite field of size q is a $(1, D, m, \mu)$-diversifying set.*

Giesbrecht and Roche [2011] also considers the case when $f \in \mathbb{C}[x]$, f is given by a numerical black box. Their proof does not apply here as the polynomials in $\mathcal{H}$ for us are not always binomials. We hope that a similar result would hold for multivariate diversification, but do not consider the question here.

Corollary 4.4 (Arnold et al. [2014], Lemma 4.1)**.** *Let bounds D, m, and μ be given as above, q a prime power, and set*

$$s = \left\lceil \log_q(2D + 1) \right\rceil \quad \text{and} \quad k = \left\lceil \log \tfrac{2}{\mu} + 2 \log m \right\rceil.$$

Then the set $\mathbb{F}_{q^s}^k$ of k-tuples from a size-s extension of the finite field with q elements is a $(1, D, m, \mu)$-diversifying set[*]*.*

[*] In this case we make the abuse of notation that each evaluation $f(\alpha x)$ is actually a k-tuple of evaluations, and the coefficients in $f(\alpha x)$ are actually k-tuples in $\mathbb{F}_{q^s}$.

We show, more generally, that vectorization may be applied to any diversifying set.

Lemma 4.5. *Let n, D, m, μ be given and $\mathcal{H}$ an unknown set of polynomials as in Definition 4.1. If $A \subset \mathsf{R}$ is a (n, D, m, μ)-diversifying set, then $A^\ell \subset \mathsf{R}^\ell$ is a (n, D, m, μ^ℓ)-diversifying set, where addition and multiplication in R^ℓ are component-wise.*

Proof. As A is a $(1, D, m, \mu)$-diversifying set, then by definition, a randomly selected row vector $\boldsymbol{\alpha} \in A^n$ satisfies $h(\boldsymbol{\alpha}) \neq 0$ for all $h \in \mathcal{H}$ with probability at least $1 - \mu$. Suppose $\boldsymbol{\beta} = (\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_\ell)$ is chosen randomly from $A^{\ell \times n}$, and note that $h(\boldsymbol{\beta}) = \big(h(\boldsymbol{\alpha}_1), \ldots, h(\boldsymbol{\alpha}_\ell)\big)$. Thus the probability that $h(\boldsymbol{\beta}) = \mathbf{0}$ is the probability that $h(\boldsymbol{\alpha}_i) = 0$ for every $i = 1, \ldots, \ell$, which is at most μ^ℓ. □

Rather than rehash the univariate diversification procedures, we refer the reader to the aforementioned results and provide the following connection which shows that univariate diversifying sets, with success probability scaled by a factor of n, become multivariate diversifying sets.

Theorem 4.6. *Let R be an integral domain and n, D, m, μ be given and $\mathcal{H}$ an unknown set of polynomials as in Definition 4.1. If $A \subset \mathsf{R}$ is a $(1, D, m, \mu/n)$-diversifying set, then A is also a (n, D, m, μ)-diversifying set.*

Proof. The proof is by induction on n. When $n = 1$, the statement holds trivially. So assume $n \geqslant 2$ and also that any $(1, d, m, \mu/(n-1))$-diversifying set is also a $(n-1, d, m, \mu)$-diversifying set.

We know that $\mathcal{H}$ is a set of n-variate polynomials, each with max degree less than D. Rewrite each h as a polynomial in x_1 with coefficients in $\mathsf{R}[x_2, \ldots, x_n]$, and define $\mathrm{lc}(h)$ to be the leading coefficient of h in terms of x_1. Each $\mathrm{lc}(h)$ is an $(n-1)$-variate polynomial with max degree less than D. Furthermore, every $\mathrm{lc}(h)$ must be nonzero since $h \neq 0$.

Now define $\mathcal{H}' = \{\mathrm{lc}(h) \mid h \in \mathcal{H}\}$. This is a set of $(n-1)$-variate polynomials with degrees less than D. Therefore, by the induction hypothesis, a random evaluation point $(\alpha_2, \ldots, \alpha_n)$ with elements chosen from A is a zero for any polynomial in $\mathcal{H}'$ with probability less than $(n-1)\mu/n$.

Now consider the set $\mathcal{H}'' = \{h(x_1, \alpha_2, \ldots, \alpha_n) \mid h \in \mathcal{H}\}$, whose leading coefficients are all nonzero with probability at least $1 - \mu/(n-1)$. $\mathcal{H}'$ is a set of at most m univariate polynomials with degrees less than D. From the original definition of A, choosing α_1 at random from A makes every $h(\alpha_1, \alpha_2, \ldots, \alpha_n)$ nonzero with probability at least $1 - \mu/n$.

Therefore a randomly-chosen point $(\alpha_1, \ldots, \alpha_n) \in A^n$ is a root of any $h \in \mathcal{H}$ only if $(\alpha_2 \ldots, \alpha_n)$ is a root of some polynomial in $\mathcal{H}'$, or α_1 is a root of some polynomial in $\mathcal{H}''$. As the probability of each of these is less than $(n-1)\mu/n$ and μ/n, respectively, the probability either occurs must be less than μ, as required. □

5 Multivariate interpolation

In this section we show how one can interpolate f using randomized substitutions and a univariate interpolation algorithm. As in (3.5), write

$$f = a_1\mathbf{x}^{\mathbf{e}_1} + \cdots + a_t\mathbf{x}^{\mathbf{e}_t} \in \mathsf{R}[x_1, \ldots, x_n].$$

with known bound $T \geqslant t$. If $T = 1$, then we simply perform n substitutions $f(z, 1, \ldots, 1)$, $\ldots$, $f(1, \ldots, 1, z)$, each of which reveals the single term and its exponent in one of the variables. No randomization is necessary in this case and the solution is trivial. Therefore for the remainder of this section we assume that $T \geqslant 2$.

5.1 Choosing multiple substitutions

The first step in interpolating f is to select ν randomized Kronecker substitutions, $\mathbf{s}_1, \ldots, \mathbf{s}_\nu \in \mathsf{R}^n$, where each $s_i = (s_{i1}, \ldots, s_{in})$. We require the $\mathbf{s}_i$ to be chosen in such a way that, with high probability, every term of f avoids collision for at least half of the substitutions $\mathbf{s}_i$.

To achieve this we first randomly select Kronecker substitutions such that any fixed term of f avoids collision for the substitution $\mathbf{s}$ with probability exceeding $3/4$.

For the bivariate case, we would choose primes $s_{i1} \in [\lambda_p, 2\lambda_p]$, and $s_{i2} \in [\lambda_q, 2\lambda_q]$, for $1 \leqslant i \leqslant \nu$, where λ_p and λ_q are determined by setting $\mu = 1/4$ in Lemma 3.3.

Applying Lemma 3.3 for the general multivariate case, we would select each integer $s_{ij} \in [0, \lambda - 1]$, where λ is the least prime greater than $4T/3$.

Given such choices of $\mathbf{s}_i$, the following lemma shows how many substitutions ν are required so that every term of f appears without collisions in at least half of them.

Lemma 5.1. *Let $f \in \mathsf{R}[x_1, \ldots, x_n]$ with max degree less than D and at most T nonzero terms. Set*

$$\nu = \max\left(4n, 8\ln\left(10T\right)\right),$$

and choose ν vectors $\mathbf{s} \in \mathbb{Z}^n$ such that, for any single $\mathbf{s}$ and any particular term in f, the probability that the term collides with another is less than $1/4$.

Then, with probability at least $9/10$, every term of f collides with no others for at least $2n$ of the substitutions.

Proof. By Hoeffding's inequality Hoeffding [1963] the probability that any fixed term of f collides in a proportion of at least $1/2$ of the substitutions is at most $\exp(-\nu/8) \leqslant 1/(10T)$. Thus the probability is at most $1/10$ that any term of f collides in more than $\nu/2 \geqslant 2n$ of the substitutions $\mathbf{s}$. □

5.2 Choosing a diversifying set

The next step is to find an appropriate diversifying set for the interpolation problem. Note the images

$$g'_i = f(z^{s_{i1}}, \ldots, z^{s_{in}}), 1 \leqslant i \leqslant \nu,$$

contain at most $T\nu$ nonzero terms in total. Every term cz^d from one of the images g'_i, $1 \leqslant i \leqslant \nu$, is the image of the (possibly empty) sum of terms of f of degree $\mathbf{e}$ satisfying $\mathbf{e}\cdot\mathbf{s}_i = d$. We take $\mathcal{H}$ to be the set of nonzero differences of all such sums. Thus, in order to obtain the appropriate diversity with some desired probability $1-\epsilon$, we require a $(n, D, \#\mathcal{H}, \epsilon)$-diversifying set. Per Theorem 4.6, it suffices that we find a $(1, D, \#\mathcal{H}, \epsilon/n)$-diversifying set A.

We randomly select $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ from A^n and then use a univariate interpolation algorithm of our choosing in order to construct the set of images

$$g_i = f(\alpha_1 z^{s_{i1}}, \ldots, \alpha_n z^{s_{in}}), \quad 1 \leqslant i \leqslant \nu$$

having the property that, with high probability, every pair of terms cz^d of g_i and cz^e of g_j $(1 \leqslant i, j \leqslant \nu)$, sharing a coefficient c, are images of the same sum of terms of f.

Lemma 5.2. *Let $f \in \mathsf{R}[x_1, \ldots, x_n]$ with max degree less than D and at most T nonzero terms. Assume ν substitution vectors are chosen according to Lemma 5.1.*

Set $m = T^2(\nu+2)^2/8$ and choose a $(n, D, m, 1/10)$-diversifying set $A \subseteq \mathsf{R}$. Then, with probability at least $4/5$, any nonzero coefficient c that appears in at least $\nu/2$ of the substitution polynomials g_i is the image of a single term in f.

Proof. From the proof of Lemma 5.1, the probability is at least $9/10$ that every term in f appears without collision in at least $\nu/2$ of the substitutions.

Assuming this is the case, there can be at most $T\nu/4$ sums of terms that collide in any image g_i, since each collision involves at least two terms, there are at least $T\nu/2$ terms that do not collide, and the total number of terms in all images is $T\nu$.

Hence the set of term differences $\mathcal{H}$ will consist of the differences of any pair in a set of $T + T\nu/4$ polynomials. The number of such pairs is less than m given in the statement of the lemma.

From the definition of A, the probability that any of these polynomials in $\mathcal{H}$ vanish on $\boldsymbol{\alpha} \in A^n$ is less than $1/10$, so the total probability that each term in f is uninvolved in collisions in at least $\nu/2$ of the images, and all distinct terms and collisions in the image polynomials g_i have distinct coefficients, is at least $(9/10)^2 > 4/5$. □

A direct consequence is that any fixed subset of terms of f must collide in fewer than half of the g_i. For every nonzero coefficient that occurs in at least $\nu/2$ of the images g_i, we know those terms with coefficient c are probably images of the same fixed term of f.

An alternate method might be to allow the $O(T\nu)$ sums of terms that appear in collisions to sometimes share the same coefficient, as long as these

coefficients are not the same as any of the T coefficients of actual terms in $f(\alpha_1 x_1, \ldots, \alpha_n x_n)$. This would reduce the m in determining the diversifying set to $T^2(\nu+2)/4$, a factor of n improvement from the bound above. The cost of such weakened diversifying sets would be that some number $\leqslant T/4$ of terms in the final recovered polynomial h are not actually terms in f. By iterating $O(\log T)$ times, such "garbage terms" could be eradicated.

5.3 Recovering the multivariate exponents

For each coefficient c that appears in at least $\nu/2$ of the images g_i, we attempt to find n linearly independent substitution vectors, call them $\mathbf{r}_1, \ldots, \mathbf{r}_n \in \{\mathbf{s}_1, \ldots, \mathbf{s}_\nu\}$, such that every substitution polynomial g_j with substitution vector $\mathbf{r}_j = (r_{j1}, \ldots, r_{jn})$, for $1 \leqslant j \leqslant n$, contains the coefficient c in a nonzero term.

In the bivariate case this is straightforward. Any 2×2 linear system formed by two substitution vectors

$$\begin{bmatrix} s_{11} & s_{12} \\ s_{21} & s_{22} \end{bmatrix}$$

must have nonzero determinant since in the bivariate case the entries are all distinct prime numbers.

The general multivariate case is more involved, as n substitution vectors may not always be linearly independent. For this case we will randomly select $2n$ vectors $\mathbf{r}_1, \ldots, \mathbf{r}_{2n} \in [0, \lambda-1]^n$, from which we will search for n linearly independent vectors. To that end we require a bound on the probability that such n independent vectors do not exist.

Lemma 5.3. *Let λ be a prime number and $f \in \mathsf{R}[x_1, \ldots, x_n]$, and suppose that a term of f avoids collision in an image $f(\mathbf{x}^{\mathbf{s}})$ for a randomly chosen $\mathbf{s} \in [0, \lambda-1]^n$ with probability at least $3/4$. Let $\mathbf{r}_1, \ldots, \mathbf{r}_{2n}$ be row vectors, chosen uniformly from $[0, \lambda-1]^n$. Given that a term of f avoids collision in the images $f(\mathbf{x}^{\mathbf{r}_i})$ for $1 \leqslant i \leqslant 2n$, then*

$$Q = \begin{bmatrix} \mathbf{r}_1 \\ \vdots \\ \mathbf{r}_{2n} \end{bmatrix}$$

has rank less than n with probability at most $(9\lambda/16)^{-n}$.

Proof. Since all entries in each $\mathbf{s}$, and thereby everything in each $\mathbf{r}$ and every element in Q, is less than λ, we can consider all these objects over the finite field $\mathbb{F}_\lambda$.

If Q has rank less than n, then $\mathbf{r}_1, \ldots \mathbf{r}_{2n}$ all lie in some dimension-$(n-1)$ subspace $W \subset \mathbb{F}_\lambda^n$. The number of distinct substitution vectors that could lie in the same subspace W is λ^{n-1}.

Each dimension-$(n-1)$ subspace $W \subset \mathbb{F}_\lambda^n$ may be specified by a nonzero vector spanning its orthogonal space, unique up to a scalar multiple. Thus

the number of such subspaces is less than λ^n, and so the number of possible $2n$-tuples comprised of substitution vectors that do not span $\mathbb{F}_\lambda^n$ is at most

$$\lambda^n(\lambda^{n-1})^{2n} = \lambda^{2n^2-n}.$$

Meanwhile, there are λ^{2n^2} possible $2n$-tuples of substitution vectors, and so the probability that such a tuple does span V is at most λ^{-n}. Furthermore, by the hypothesis, the probability that a term of f avoids collision for each substitution $\mathbf{r}_i$ is $(3/4)^{2n}$, and thus the conditional probability that Q is not full rank given that a fixed term of f avoids collision for each $\mathbf{r}_i$ is at most $\lambda^{-n}/(3/4)^{2n} = (9\lambda/16)^{-n}$. □

Given such a high probability of each term producing a rank-n system of substitution vectors, it is a simple matter to show that with high probability *every* term of f admits some such rank-n linear system of substitutions without collisions.

Corollary 5.4. *Let $f \in \mathsf{R}[x_1, \ldots, x_n]$ as above, and set ν according to Lemma 5.1, $\boldsymbol{\alpha}$ according to Lemma 5.2, and $\lambda \geqslant 3$ according to Lemma 3.3. With probability at least $2/3$, for every term in f, there exists a rank-n set of substitution vectors $\mathbf{r}_1, \ldots, \mathbf{r}_n$, such that the given term of f does not collide in any of the substitutions $g = f(\alpha_1 z^{r_{i1}}, \ldots \alpha_n z^{r_{in}})$, for $1 \leqslant i \leqslant n$.*

Proof. As we have discussed, the case $T = 1$ is trivial and when $n = 2$ we choose primes for the vectors $\mathbf{s}$ and the 2×2 linear systems always have full rank.

So assume $T \geqslant 2$ and $n \geqslant 3$. We know from Lemma 5.3 that the probability of a single term *not* admitting a rank-n system of substitution vectors is less than $(9\lambda/16)^{-n}$, so the probability that any term does not have a rank-n system of non-colliding substitution vectors is less than $T/(9\lambda/16)^n$.

Since λ is chosen as the least odd prime greater than $4T/3$, we see that $T/(9\lambda/16)^n \geqslant \frac{4}{3}(9\lambda/16)^{-n+1}$. And because $\lambda \geqslant 3$ and $n \geqslant 3$, $\frac{4}{3}(9\lambda/16)^{-n+1} \leqslant \frac{4}{3}(27/16)^{-2} < 1/6$.

Combining this with the probability bound from Lemma 5.2, the overall success probability is at least $\frac{5}{6} \cdot \frac{4}{5} = \frac{2}{3}$. □

From the set of $2n$ vectors $\mathbf{r}_1, \ldots, \mathbf{r}_{2n}$ we can find n linearly independent vectors by inspection of the LU factorization of the matrix whose $2n$ rows are the $\mathbf{r}_j$. By Bunch and Hopcroft [1974], we can do this in $\widetilde{O}(n^\omega)$ operations in $\mathbb{F}_\lambda$, for a bit cost of $\widetilde{O}(n^\omega \log T)$. We suppose by reordering of the $\mathbf{r}_j$ that $\mathbf{r}_1, \ldots, \mathbf{r}_n$ are our n linearly independent vectors.

Then, if d_j is the exponent of the term with coefficient c appearing in $f(\alpha_1 z^{r_{j1}}, \ldots, \alpha_n z^{r_{jn}})$, we may find the degree $\mathbf{e}$ of the term with coefficient c in the diversified multivariate polynomial $f(\alpha_1 x_1, \ldots, \alpha_n x_n)$ by way of the linear system

$$\begin{bmatrix} \mathbf{r}_1 \\ \vdots \\ \mathbf{r}_n \end{bmatrix} \begin{bmatrix} e_1 \\ \vdots \\ e_n \end{bmatrix} = \begin{bmatrix} d_1 \\ \vdots \\ d_n \end{bmatrix}.$$

We construct and solve such a linear system for every term of f, giving us the polynomial

$$g = f(\alpha_1 x_1, \ldots, \alpha_n x_n),$$

from which we easily obtain f as $f = g(\alpha_1^{-1} x_1, \ldots, \alpha_n^{-n} x_n)$. Procedure Interpolate describes the approach laid out in sections 5.1-5.3.

Procedure Interpolate(f, n, T, D)

Input: Bounds D, T and a black box for evaluating $f \in \mathsf{R}[x_1, \ldots, x_n]$, an unknown polynomial with partial degrees less than D and at most T nonzero terms.

Result: We construct f with probability at least 2/3.

`// Choose substitution vectors`
$\nu \leftarrow \max(4n, 8\ln(10T))$
if $n = 2$ **then**
  $\mathbf{s}_1, \ldots, \mathbf{s}_\nu \leftarrow$ prime vectors chosen by Lemma 3.1
else
  $\mathbf{s}_1 \ldots, \mathbf{s}_\nu \leftarrow$ integer vectors chosen by Lemma 3.3

`// Diversify`
$m \leftarrow T^2(\nu+2)^2/8$
$A \leftarrow$ a $(n, D, m, 1/10)$-diversifying subset of R
$\boldsymbol{\alpha} \leftarrow$ an element of A^n chosen uniformly at random

`// Build images of` f
for $i = 1, \ldots, \nu$ **do**
  $g_i \leftarrow f(\alpha_1 z^{s_{i1}}, \ldots, \alpha_n z^{s_{in}})$ via univariate interpolation

`// Reconstruct` f `from its images`
$g \leftarrow 0$
foreach *coefficient* $c \neq 0$ *appearing in any* g_i **do**
  $S \leftarrow \{\mathbf{s}_i \mid c \text{ is a coefficient in } g_i\}$
  if $\#S < \nu/2$ **then continue**
  $\mathbf{r}_1, \ldots, \mathbf{r}_n \leftarrow$ linearly independent vectors chosen from the first $2n$ vectors in S
  $d_1, \ldots, d_n \leftarrow$ degrees of terms with coefficient c under substitutions $\mathbf{r}_1, \ldots, \mathbf{r}_n$
  $R \leftarrow (\mathbf{r}_1 \ \ldots \ \mathbf{r}_n)^\top$
  $\mathbf{d} \leftarrow (d_1 \ldots d_n)^\top$
  $\mathbf{e} \leftarrow R^{-1}\mathbf{d}$
  $g \leftarrow g + c\mathbf{x}^{\mathbf{e}}$
return $g(\alpha_1^{-1} x_1, \ldots, \alpha_n^{-1} x_n)$

5.4 Cost analysis

We can now state the tangible benefit of our new Kronecker substitution technique. The diversifying sets are included in these theorems even though we do not actually count the cost of possibly extending the ring R to include such sets.

Theorem 5.5. *For given bounds D_x, D_y, T and an unknown polynomial $f \in \mathsf{R}[x,y]$ with $\deg_x f < D_x$, $\deg_y f < D_y$, and $\#f \leqslant T$, Procedure Interpolate succeeds in finding f with probability at least $2/3$ and requires*

- *$O(\log T)$ calls to univariate interpolation with T nonzero terms and degree $O\left(\sqrt{T}\sqrt{D_x D_y}\log(D_x D_y)\right)$,*
- *A $\left(2, D, O\left(T^2\log^2 T\right), 1/10\right)$-diversifying set in R, and*
- *$\widetilde{O}\left(T\log D + \log^2 D\right)$ additional bit operations, where $D = \max(D_x, D_y)$.*

Proof. Since $n = 2$, we have $\nu \in O(\log T)$. This is the number of calls to the univariate interpolation algorithm, and the degree bound comes from Corollary 3.2. The size of $\mathcal{H}$ in the diversifying set comes from the fact that $m \in O\left(T^2\log^2 T\right)$.

Two steps dominate the bit complexity. First, we must choose 2ν primes in $[\lambda, 2\lambda]$. This can be accomplished via $O(\nu)$ applications of the Miller-Rabin primality test, performing $O(\log\log T)$ trials each time to ensure a negligible probability of error. The cost of these tests is $\widetilde{O}\left(\log T\log^2\lambda\right)$, which is $\widetilde{O}\left(\log T\log^2 D\right)$.

The other dominating step in bit complexity is simply the cost of computing with the $T\nu$ exponents in images g_i and T exponent vectors in the final result. There are $\widetilde{O}(T)$ such exponents, each with $O(\log D)$ bits, for a total bit cost of $\widetilde{O}(T\log D)$. □

Theorem 5.6. *For given bounds D, T and an unknown polynomial $f \in \mathsf{R}[x_1, \ldots, x_n]$ with $\max\deg g < D$ and $\#f \leqslant T$, Procedure Interpolate succeeds in finding f with probability at least $2/3$ and requires*

- *$O(n + \log T)$ calls to univariate interpolation with T nonzero terms and degree $O(TD)$,*
- *A $\left(2, D, \widetilde{O}\left(T^2n^2\right), 1/10\right)$-diversifying set in R, and*
- *$\widetilde{O}\left(n^\omega T + nT\log D\right)$ additional bit operations, where $2 < \omega < 3$ is the exponent of matrix multiplication.*

Proof. The analysis of the first two parts is the same as in the bivariate case.

For the bit complexity, we do not have to worry about primality testing here. However, the size of all exponents in the polynomials becomes $\widetilde{O}(nT\log D)$, and the cost of performing each LU factorization on a $(2n)\times n$ matrix is $O(n^\omega)$ operations on integers with $O(\log T)$ bits. As T such LU factorizations are required, the total bit cost of the linear algebra is $\widetilde{O}(n^\omega T)$. □

5.5 Interpolating f with arbitrarily high probability

Interpolating f entails the probabilistic steps of (1) selecting a set of randomized substitutions that produce few collisions and (2) selecting $\boldsymbol{\alpha}$ from a diversifying

set A such that all term sums in all images have distinct coefficients in those images. The $n \geqslant 3$ case has in addition the probabilistic step (3) of guaranteeing that we can construct a full-rank linear system in order to solve for every exponent of f. The probability of failure in each of these steps has been controlled above so that the overall success probability is at least 2/3.

If a higher success probability, say $1 - \epsilon$, is desired, we simply run the interpolation algorithm described in sections 5.1–5.3 with some ℓ times. Again using Hoeffding's inequality, the probability that the algorithm fails at least $\ell/2$ times is at most $\exp(-2\ell(1/6)^2) = \exp(-\ell/18)$. Thus, if we wish to discover f with probability $1 - \epsilon$, we merely run the algorithm as suggested some $\ell = \lceil 18 \ln \frac{1}{\epsilon} \rceil$ times, and select the the polynomial f that is returned a majority of the time. With probability at least $1 - \epsilon$, such an f exists and is in fact the correct answer.

6 Perspective

We have presented a new randomization that maps a multivariate polynomial to a univariate polynomial with (mostly) the same terms. This improves on the usual Kronecker map by reducing the degree of the univariate image when the polynomial is known to be sparse. We have also shown how a small number of such images can be combined to recover the original terms of the unknown multivariate polynomial.

There are numerous questions raised by this result. Perhaps foremost is whether there is any practical gain in any particular application by using this approach. We know that the randomized Kronecker substitution will result in smaller degrees than the usual Kronecker substitution whenever the polynomial is sufficiently large and sufficiently sparse, so in principle the applications should include any of the numerous results on sparse polynomials that use a Kronecker substitution to accommodate multivariate polynomials.

Unfortunately, in practice, the situation is not so clear. Many of the aforementioned results that rely on a Kronecker substitution either do not have a widely-available implementation, or do not usually involve sparse polynomials. However, for the particular applications of sparse GCD and sparse multivariate multiplication, there is considerable promise particularly in the case of bivariate polynomials with degree greater than 1000 or so and sparsity between D and D^2. An efficient implementation comparison in these situations would be useful and interesting, and we are working in that direction.

There are also questions of theoretical interest. For one, we would like to know how far off the bounds on the size of primes from Lemmata 3.1 and 3.3 are compared to what is really necessary to avoid collisions.

An important question is whether our current results are optimal in any sense. In the bivariate case, when $D_x = D_y$ our result gives $p, q \in \widetilde{O}\left(\sqrt{T}\right)$, which is optimal in terms of T. That is because T could be as large as $\Theta(D^n)$, and therefore any monomial substitution exponent less than $\Omega(T^{(n-1)/n})$ would by necessity have more than a constant fraction of collisions. However our result

for $n \geqslant 2$ gives each $p_i \in \widetilde{O}\left(T + n^2 + \log^2 D\right)$, which in terms of T is off by a factor of $T^{1/n}$ from the optimal. It may be possible to improve these bounds simply with a better analysis, or with a different kind of randomized monomial substitution. In either case, it is clear that, at least for $n \geqslant 3$ and in particular for $n = 3$, it should be possible to improve on the results here and achieve univariate reduced polynomials with even lower degree.

Another interesting question would be whether some of this randomization can be avoided. Here we have two randomizations, the diversification and the (multiple) randomized Kronecker substitutions. And this is besides any randomization that might occur in the underlying univariate algorithm! It seems plausible that, for example in the application of multivariate multiplication, the known aspects of the monomial structure might be used to make some choices less random and more "intelligent". However, we do not yet know any reasonable way to accomplish this.

7 Acknowledgements

We wish to thank Zeev Dvir for pointing out the previous work of Klivans and Spielman, and the reviewers for their helpful comments. We also thank the organizers of the SIAM AG13 meeting for the opportunity to discuss preliminary work on this topic. The second author is supported by the National Science Foundation, award #1319994.

A version of this paper will appear at ISSAC 2014 in Kobe, Japan.

References

Andrew Arnold, Mark Giesbrecht, and Daniel S. Roche. Sparse interpolation over finite fields via low-order roots of unity. In *ISSAC 2014 (to appear)*, 2014. URL http://arxiv.org/abs/1401.4744. Referenced on pages 8 and 9.

Martín Avendaño, Teresa Krick, and Ariel Pacetti. Newton-Hensel interpolation lifting. *Found. Comput. Math.*, 6(1):81–120, 2006. doi: 10.1007/s10208-005-0172-3. Referenced on page 3.

Michael Ben-Or and Prasoon Tiwari. A deterministic algorithm for sparse multivariate polynomial interpolation. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, STOC '88, pages 301–309, New York, NY, USA, 1988. ACM. doi: 10.1145/62212.62241. Referenced on page 3.

Jérémy Berthomieu and Grégoire Lecerf. Reduction of bivariate polynomials from convex-dense to dense, with application to factorizations. *Math. Comp.*, 81:1799–1821, 2012. doi: 10.1090/S0025-5718-2011-02562-7. Referenced on page 3.

James R. Bunch and John E. Hopcroft. Triangular factorization and inversion by fast matrix multiplication. *Mathematics of Computation*, 28(125):231–236, 1974. URL http://www.jstor.org/stable/2005828. Referenced on page 14.

Emmanuel J. Candés, Justin K. Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006. doi: 10.1002/cpa.20124. Referenced on page 3.

Matthew T. Comer, Erich L. Kaltofen, and Clément Pernet. Sparse polynomial interpolation and Berlekamp/Massey algorithms that correct outlier errors in input values. In *Proceedings of the 37th International Symposium on Symbolic and Algebraic Computation*, ISSAC '12, pages 138–145, New York, NY, USA, 2012. ACM. doi: 10.1145/2442829.2442852. Referenced on page 3.

Mark Giesbrecht and Daniel S. Roche. Diversification improves interpolation. In *Proceedings of the 36th international symposium on Symbolic and algebraic computation*, ISSAC '11, pages 123–130, New York, NY, USA, 2011. ACM. doi: 10.1145/1993886.1993909. Referenced on pages 8 and 9.

Mark Giesbrecht, George Labahn, and Wen-shin Lee. Symbolic-numeric sparse interpolation of multivariate polynomials. *Journal of Symbolic Computation*, 44(8):943 – 959, 2009. doi: 10.1016/j.jsc.2008.11.003. Referenced on page 3.

Dima Yu. Grigoriev, Marek Karpinski, and Michael F. Singer. Fast parallel algorithms for sparse multivariate polynomial interpolation over finite fields. *SIAM Journal on Computing*, 19(6):1059–1063, 1990. doi: 10.1137/0219073. Referenced on page 3.

David Harvey. Faster polynomial multiplication via multipoint Kronecker substitution. *Journal of Symbolic Computation*, 44(10):1502–1510, 2009. doi: 10.1016/j.jsc.2009.05.004. Referenced on pages 3 and 5.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, 58:13–30, 1963. URL http://www.jstor.org/stable/2282952. Referenced on page 11.

Joris van der Hoeven and Grégoire Lecerf. Sparse polynomial interpolation in practice. April 2014. URL http://hal.archives-ouvertes.fr/hal-00980366. Referenced on page 3.

Ming-Deh A Huang and Ashwin J Rao. Interpolation of sparse multivariate polynomials over large finite fields with applications. *Journal of Algorithms*, 33(2):204–228, 1999. doi: 10.1006/jagm.1999.1045. Referenced on page 3.

Seyed Mohammad Mahdi Javadi and Michael Monagan. On factorization of multivariate polynomials over algebraic number and function fields. In *Proceedings of the 2009 international symposium on Symbolic and algebraic computation*, ISSAC '09, pages 199–206, New York, NY, USA, 2009. ACM. doi: 10.1145/1576702.1576731. Referenced on page 3.

Seyed Mohammad Mahdi Javadi and Michael Monagan. Parallel sparse polynomial interpolation over finite fields. In *Proceedings of the 4th International Workshop on Parallel and Symbolic Computation*, PASCO '10, pages 160–168, New York, NY, USA, 2010. ACM. doi: 10.1145/1837210.1837233. Referenced on page 3.

Erich Kaltofen and Wen-shin Lee. Early termination in sparse interpolation algorithms. *Journal of Symbolic Computation*, 36(3-4):365–400, 2003. doi: 10.1016/S0747-7171(03)00088-9. ISSAC 2002. Referenced on page 3.

Erich L. Kaltofen. Fifteen years after DSC and WLSS2: What parallel computations I do today [invited lecture at PASCO 2010]. In *Proceedings of the 4th International Workshop on Parallel and Symbolic Computation*, PASCO '10, pages 10–17, New York, NY, USA, 2010. ACM. doi: 10.1145/1837210.1837213. Referenced on page 3.

Adam R. Klivans and Daniel Spielman. Randomness efficient identity testing of multivariate polynomials. In *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing*, STOC '01, pages 216–223, New York, NY, USA, 2001. ACM. doi: 10.1145/380752.380801. Referenced on page 3.

Leopold Kronecker. Grundzüge einer arithmetischen Theorie der algebraischen Grössen. *Journal für die reine und angewandte Mathematik*, 92:1–122, 1882. Referenced on page 3.

Yishay Mansour. Randomized interpolation and approximation of sparse polynomials. *SIAM Journal on Computing*, 24(2):357–368, 1995. doi: 10.1137/S0097539792239291. Referenced on page 3.

J. Barkley Rosser and Lowell Schoenfeld. Approximate formulas for some functions of prime numbers. *Ill. J. Math.*, 6:64–94, 1962. URL http://projecteuclid.org/euclid.ijm/1255631807. Referenced on page 5.

Arnold Schönhage. Asymptotically fast algorithms for the numerical multiplication and division of polynomials with complex coefficients. In Jacques Calmet, editor, *Computer Algebra*, volume 144 of *Lecture Notes in Computer Science*, pages 3–15. Springer Berlin Heidelberg, 1982. doi: 10.1007/3-540-11607-9_1. Referenced on pages 3 and 5.

Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5(3-4):207–388, 2010. doi: 10.1561/0400000039. Referenced on page 3.

Richard Zippel. Probabilistic algorithms for sparse polynomials. In Edward Ng, editor, *Symbolic and Algebraic Computation*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226. Springer Berlin / Heidelberg, 1979. doi: 10.1007/3-540-09519-5_73. Referenced on page 3.

Richard Zippel. Interpolating polynomials from their values. *Journal of Symbolic Computation*, 9(3):375–403, 1990. doi: 10.1016/S0747-7171(08)80018-1. Computational algebraic complexity editorial. Referenced on page 3.